Stochastic dynamics on slow manifolds

George W A Constable[1], Alan J McKane[1], Tim Rogers[2]

[1] Theoretical Physics, School of Physics and Astronomy, The University of Manchester, Manchester M13 9PL, UK
[2] Department of Mathematical Sciences, University of Bath, Claverton Down, Bath BA2 7AY, UK

Abstract. The theory of slow manifolds is an important tool in the study of deterministic dynamical systems, giving a practical method by which to reduce the number of relevant degrees of freedom in a model, thereby often resulting in a considerable simplification. In this article we demonstrate how the same basic methodology may also be applied to stochastic dynamical systems, by examining the behaviour of trajectories conditioned on the event that they do not depart the slow manifold. We apply the method to two models: one from ecology and one from epidemiology, achieving a reduction in model dimension and illustrating the high quality of the analytical approximations.

1. Introduction

It is possibly only a slight exaggeration to say that of all the mathematical models we can dream of, there are only two kinds which are straightforward to solve: those which are linear, and those which are one-dimensional. This aphorism holds equally for stochastic dynamical systems as it does for their deterministic counterparts. Much of applied mathematics and theoretical physics is devoted to the delicate art of taking high-dimensional or non-linear problems of interest and finding appropriate approximation schema with which to reduce their apparent difficultly.

The theory of 'slow' or 'centre' manifolds is an example of just such a scheme and one which is well developed for deterministic dynamical systems [1]. In many models of interest there exists a separation of time scales between some quantities which relax very quickly to an essentially static value, and others which change more slowly and can be sensitive to perturbations. The term 'slow manifold' describes the space in which these slower quantities vary, after any fast initial transient has died out. Restricting attention to this space offers an effective method by which to remove the so-called fast degrees of freedom, thereby reducing the dimension and simplifying the system. In this article we will show how the principles of slow manifold theory from deterministic systems may be put to good use in stochastic systems as a tool to separate time scales and reduce the dimensionality of the model.

The goal of removing fast degrees of freedom from stochastic systems has received significant attention and a variety of approximation methods have previously been proposed [2–11], though a simple and generally applicable theory is yet to emerge. The existing literature on the subject can be coarsely split according to the framework within which the stochastic system is represented. Some authors have focused on master or Fokker-Planck type equations which govern the time evolution of the probability distribution of system states [12]. Within this formalism, a reduction of dimension can be achieved through the application of projection operators, as illustrated by Gardiner [4], and developed by others [5, 6]. Unfortunately, the details of the method can be somewhat cumbersome, especially in cases where the fast degrees of freedom are not parallel to the variables with which the system is described. The alternative formalism of stochastic differential equations (SDEs) provides a perhaps more intuitive approach. The equations describe the evolution of trajectories themselves and bear a useful resemblance to ordinary differential equations, which aids physical reasoning. For this reason we choose the SDE formulation as the basis for our work. However, it should be pointed out that such systems cannot be correctly interpreted without specifying the choice of stochastic calculus, and that certain manipulations are not entirely straightforward [13].

Perhaps the simplest implementation of time scale separation in the SDE setting is the treatment of 'direct adiabatic elimination' presented in [7]. It will be instructive to review the salient points. The procedure mirrors closely that of slow manifold theory: the variables associated with the fast dynamics are assumed to be stationary, from which

a function describing the slow manifold may be determined. If the fast variables are subject to noise, then the procedure will yield an expression for the slow manifold which contains noise variables; unfortunately this can lead to ill-defined terms [7]. The method can, however, be usefully applied in systems where the entire system is linear in the fast variables. The more rigorously derived and well-known Haken slaving principle [2,3,14], and several other related methods [8–10] are developed along similar lines of reasoning and suffer from the same complications which arise from specifying the slow manifold in a stochastic sense. A notable exception is [15], which deals with a particular model in which there is a true centre manifold (that is, a surface on which there is no deterministic flow), and applies a novel method in which stochastic perturbations away from the manifold are assumed to instantaneously relax along the deterministic trajectories. We have applied a similar procedure in several previous works [11,16,17], and it is this idea which will form the basis of the method we develop here.

The core of our approach is to examine the behaviour of a stochastic system in the SDE framework under the condition that its trajectories are *confined* to the slow manifold of the deterministic version of the system. Because we use a static description of the slow manifold, the method is applicable to a broader range of systems than the direct elimination procedure or the Haken slaving principle. Moreover, the procedure is mathematically explicit and straightforward to apply. One also gains a sense of physical intuition as to the behaviour of the system, which is arguably not present in the master equation setting.

We develop and demonstrate the method with the aid of a pair of motivating examples, which we select from the burgeoning field of research into the effects of demographic noise in individual-based models (IBMs). Models of this type have recently become very popular in several areas of physics and applied mathematics as tools to study complex emergent phenomena. In the limit of large system size, the mesoscopic behaviour of a given IBM is well described by an SDE system derived from the rules of the model (see [17] and the appendices of this paper for details). It is to these SDEs that we will apply our method.

We will also show how the slow manifold approximation can be effectively combined with other stochastic approximation techniques. The linear noise approximation (LNA, or van Kampen expansion [18]) has found favour amongst theoreticians studying IBMs. In a nutshell, the LNA provides a mesoscopic description of the system in terms of a linear SDE. The price paid for this simplification is that the theory only applies in the neighbourhood of an attractive fixed point of the noise-free version of the model. In the context of the LNA, slow manifolds can be a malign presence. If some eigenvalues of the approximate linear SDE are close to zero, then a small stochastic fluctuation in the direction of the corresponding eigenvector can carry the system very far from the steady state into regions in which the true non-linear nature of the model is important. Moreover, a large separation between eigenvalues can in some situations lead to the numerical evaluation of theoretical solutions becoming ill-conditioned. Both of these effects can lead to a very poor agreement between stochastic simulations and the LNA

theory.

In the next section we develop our method with the aid of a simple illustrative example from ecology, before providing a general formulation. The results from this example demonstrate the success of the approximation scheme even in regimes where the fixed point is weakly unstable; this addresses the first difficulty with slow manifolds and the LNA identified above. In section three we go on to apply the general formulation of the method to an epidemiological model with seasonal forcing. This model has been identified as suffering from the technical numerical difficulties associated with a large separation between eigenvalues [19]. We show how our method may be used in tandem with the LNA to provide a very good approximation to results coming from stochastic simulations.

2. Method

2.1. Motivation

We begin by recapping the basics of slow manifolds in deterministic systems. Consider the ordinary differential equation

$$\frac{d\boldsymbol{x}}{dt} = \boldsymbol{A}(\boldsymbol{x}), \tag{1}$$

where $\boldsymbol{x}$ is an n-dimensional vector describing the state of the system, and $\boldsymbol{A}$ is an n-dimensional vector-valued function of $\boldsymbol{x}$. As is well known, the behaviour of the system in the neighbourhood of a fixed point $\boldsymbol{x}^*$ is described by the linearization of $\boldsymbol{A}$ around that point. Define the Jacobian matrix J with entries

$$J_{ij} = \frac{\partial}{\partial x_j} A_i(\boldsymbol{x})\Big|_{\boldsymbol{x}=\boldsymbol{x}^*}. \tag{2}$$

Then for $\boldsymbol{x}$ close to $\boldsymbol{x}^*$, the time evolution of $\boldsymbol{\xi} = \boldsymbol{x} - \boldsymbol{x}^*$ obeys

$$\frac{d\boldsymbol{\xi}}{dt} = J\boldsymbol{\xi}. \tag{3}$$

Further insight is gained by considering the eigenvalues and eigenvectors of J. From (3), we learn that if $\boldsymbol{v}$ is a left eigenvector of J with eigenvalue λ, then errors in the direction of $\boldsymbol{v}$ will grow exponentially if $\lambda > 0$ and shrink exponentially if $\lambda < 0$. If, on the other hand, $\lambda = 0$ then we do not know what effect perturbations in the direction of $\boldsymbol{v}$ will have on the long-term behaviour of the system. To answer this question would require a more detailed non-linear analysis, which is likely to be very difficult in a general system with several degrees of freedom. Slow manifold theory offers a way to make progress in this case by reducing the dimension of the model.

The basic observation behind the theory is as follows: since perturbations in the direction of stable/unstable eigenvectors will shrink/grow exponentially, the only trajectories whose behaviour is in question are those which are tangential to the span of the eigenvectors with eigenvalue zero. Very often, this set of eigenvectors has many fewer members than there are degrees of freedom in the original system, and thus restricting attention to this subspace achieves a considerable reduction in dimensionality.

Slow manifolds are also of great practical use when no eigenvalues are precisely zero, but there is a separation of time-scales. For example, suppose a stable fixed point $\boldsymbol{x}^*$ has associated eigenvalues satisfying $\mathrm{Re}[\lambda_1] < \ldots < \mathrm{Re}[\lambda_m] \ll \mathrm{Re}[\lambda_{m+1}] < \ldots < 0$. Perturbations in the direction of eigenvectors $\boldsymbol{v}_1, \ldots, \boldsymbol{v}_m$ will decay extremely rapidly in comparison with those in the directions of $\boldsymbol{v}_{m+1}, \ldots, \boldsymbol{v}_n$. For practical purposes, the 'slow' manifold of trajectories tangent to these less stable eigenvectors defines an $(n - m)$-dimensional system which will provide a good qualitative approximation to the behaviour of the larger system, as perturbations away from this manifold will very quickly collapse.

Our goal in this article is to apply the basic ideas of slow manifolds to stochastic systems. Our starting point will be the SDE

$$\frac{d\boldsymbol{x}}{dt} = \boldsymbol{A}(\boldsymbol{x}) + \boldsymbol{\eta}(t), \tag{4}$$

where $\boldsymbol{x}$ and $\boldsymbol{A}$ are as in (1), and $\boldsymbol{\eta}$ is a vector of Gaussian white-noise variables with zero mean and correlations

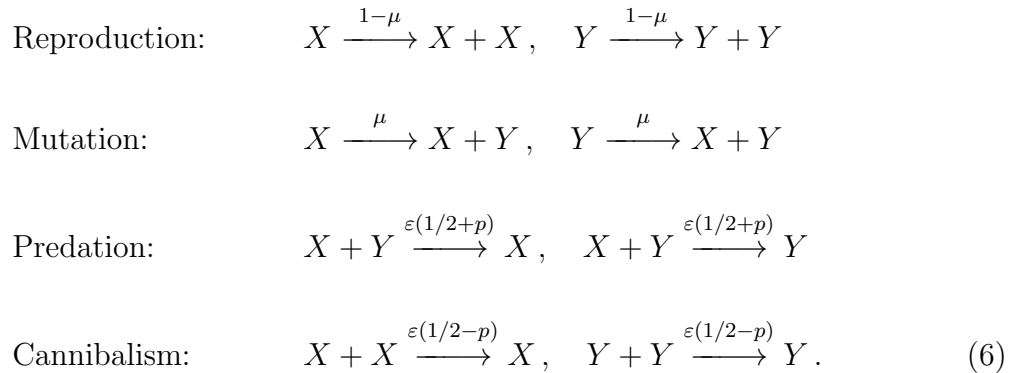$$\left\langle \eta_i(t)\eta_j(t') \right\rangle = \varepsilon\, \delta(t - t') B_{ij}(\boldsymbol{x}). \tag{5}$$

Here $\langle \cdots \rangle$ denotes averaging over the noise, ε is a small parameter governing the strength of the noise, B is a matrix-valued function of the system state, and the SDE (4) is to be interpreted in the Itō sense. The exact form of $\boldsymbol{A}$ and B will be determined by the application.

We are interested in the case where the deterministic ($\varepsilon = 0$) system exhibits a slow manifold. How will the stochastic system behave if we confine its trajectories to this slow manifold? We develop the theory with the aid of a specific example.

2.2. Illustrative example

To illustrate our method, we explore the behaviour of a simple ecological model of two interacting populations, labelled X and Y. Individuals of both populations reproduce with rate one, and there is a small probability μ of the offspring mutating from one type to the other. The organisms also prey on each other with rate ε and a slight preference p for prey of the opposite type. The model may be written in the traditional notation of chemical reaction systems:

$$\text{Reproduction:} \qquad X \xrightarrow{1-\mu} X + X, \quad Y \xrightarrow{1-\mu} Y + Y$$

$$\text{Mutation:} \qquad X \xrightarrow{\mu} X + Y, \quad Y \xrightarrow{\mu} X + Y$$

$$\text{Predation:} \qquad X + Y \xrightarrow{\varepsilon(1/2+p)} X, \quad X + Y \xrightarrow{\varepsilon(1/2+p)} Y$$

$$\text{Cannibalism:} \qquad X + X \xrightarrow{\varepsilon(1/2-p)} X, \quad Y + Y \xrightarrow{\varepsilon(1/2-p)} Y. \tag{6}$$

Here arrows denote possible reactions and the values above them are the rate constants. Writing n_X and n_Y for the number of individuals in each population, the model may be mathematically formulated as a master equation describing the time evolution of the probability distribution $P(n_X, n_Y)$; see (A.2) in Appendix A. Stochastic simulations of the model can be performed efficiently using the Gillespie algorithm [20].

When the predation rate ε is small, the population may grow very large. Performing an expansion of the master equation in the limit of small ε yields an effective description of the system in terms of an SDE for the scaled variables $x = \varepsilon n_X$ and $y = \varepsilon n_Y$. We give the details in Appendix A. For the present model, we find the following pair of equations:

$$\frac{dx}{dt} = x - \mu(x - y) - x\left(\frac{1}{2}(x + y) - p(x - y)\right) + \eta_x(t),$$

$$\frac{dy}{dt} = y + \mu(x - y) - y\left(\frac{1}{2}(x + y) + p(x - y)\right) + \eta_y(t),$$

(7)

where η_x and η_y have the correlation structure specified in (5), with

$$B = \begin{pmatrix} x + \frac{1}{2}x(x + y) - (px + \mu)(x - y) & 0 \\ 0 & y + \frac{1}{2}y(x + y) + (py + \mu)(x - y) \end{pmatrix}.$$

(8)

We begin by examining the deterministic system found by putting $\varepsilon = 0$. There is a trivial fixed point at $x^* = 0$, $y^* = 0$, representing the extinct state, which is always unstable. There is a second fixed point at $x^* = 1$, $y^* = 1$, representing equal coexistence of the two populations. This state is stable when $p < \mu$. If p is raised above μ, a supercritical pitchfork bifurcation occurs, with the equal coexistence fixed point becoming unstable and giving rise to a symmetric pair of stable fixed points in which one species dominates the other. The new fixed points have coordinates

$$x^* = \frac{1 - 2\mu \pm \sqrt{(1 - \mu/p)(1 - 2\mu)}}{1 - 2p}, \quad y^* = \frac{1 - 2\mu \mp \sqrt{(1 - \mu/p)(1 - 2\mu)}}{1 - 2p}.$$

(9)

We are interested in examining the effect of noise near this transition.

The eigenvalues of the Jacobian at the coexistence state are -1 and $\lambda \equiv 2(p - \mu)$, with corresponding eigenvectors $(1, 1)$ and $(1, -1)$. If $|\lambda| \ll 1$ then we have a slow manifold in the direction of $(x - y)$, meaning that perturbations to the balance of populations evolve very slowly. Formally, the slow manifold is defined by the collection of trajectories which are tangent to the slow eigenvector at the fixed point, although in practice there is unlikely to be a closed analytic expression for this surface. To make progress we approximate the slow manifold by the surface on which the rate of change in the direction of the fast eigenvector is zero (known as the nullcline). In the present model, the nullcline of the fast direction $(x + y)$ is given by $dx/dt + dy/dt = 0$ which yields the hyperbola

$$(x + y) - \frac{1}{2}(x + y)^2 + p(x - y)^2 = 0.$$
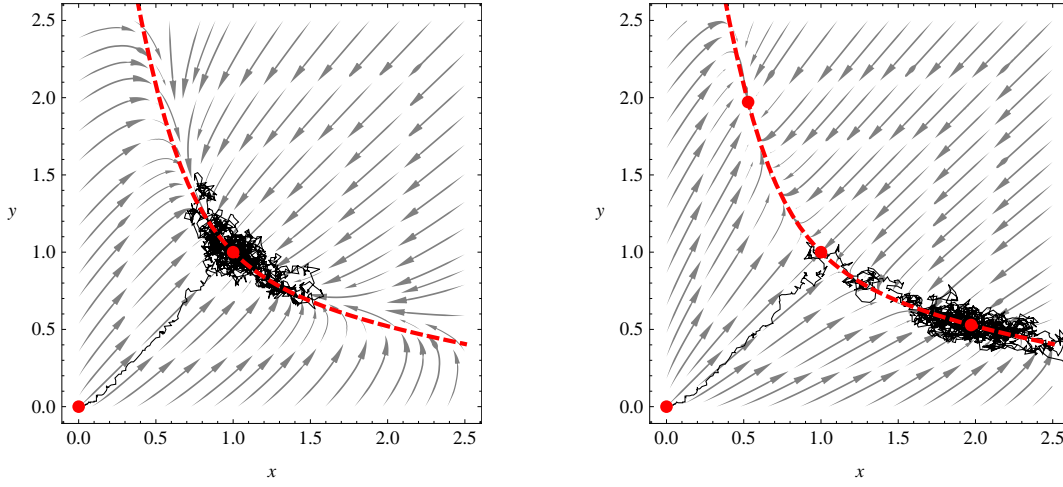
(10)

Figure 1. These plots show the behaviour of our example illustrative ecological model either side of the pitchfork bifurcation. The fixed points are shown as red circles, and the dashed red line is the nullcline $dx/dt + dy/dt = 0$, given in (10). The grey arrows show trajectories of the deterministic system, while the black line traces out the trajectory of a single (short) stochastic simulation of the individual-based model, starting close to the origin. The parameters are $\varepsilon = 0.005$ and $p = 0.3$ in both plots, while $\mu = 0.35$ on the left and $\mu = 0.25$ on the right.

The two plots in figure 1 capture the typical behaviour of the model for parameters either side of the transition.

The SDE system (7) is two-dimensional, non-linear and has noise correlations which depend on the state of the system. These factors combine to make the theoretical analysis of the model very difficult. The situation is not hopeless, however, as it is clearly visible in figure 1 that the system state does not typically stray very far from one-dimensional subspace defined by the nullcline of the fast variable $(x+y)$. We intend to exploit this fact to produce an 'effective' one-dimensional description of the model. The plan of attack is as follows: first we will make a coordinate transform to separate the fast and slow variables; then we will examine the behaviour of the slow variable under the assumption that the fast variable relaxes instantaneously to its nullcline value.

We introduce $w = x + y$ and $z = x - y$, so that

$$\begin{pmatrix} w \\ z \end{pmatrix} = \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} \equiv V \begin{pmatrix} x \\ y \end{pmatrix}. \tag{11}$$

In the new coordinates the nullcline is described by the equation $w - w^2/2 + pz^2 = 0$, and (7) becomes

$$\frac{dw}{dt} = w - \frac{1}{2}w^2 + pz^2 + \eta_w(t),$$

$$\frac{dz}{dt} = z\left(1 - 2\mu - \left(\frac{1}{2} - p\right)w\right) + \eta_z(t). \tag{12}$$

To determine the correlation structure of the new noise variables η_w and η_z, we apply a general result on Gaussian random variables. Suppose that a vector of Gaussian random variables $\boldsymbol{\eta}$ has correlation matrix B, and that $\boldsymbol{\zeta} = V\boldsymbol{\eta}$ for some matrix V. What is the correlation matrix of $\boldsymbol{\zeta}$? Well,

$$\langle \zeta_i \zeta_j \rangle = \left\langle \sum_{k,l} V_{ik}\eta_k V_{jl}\eta_l \right\rangle = \sum_{k,l} V_{ik}\langle \eta_k \eta_l \rangle V_{lj}^T = \tilde{B}_{ij}\,, \tag{13}$$

where $\tilde{B} = VBV^T$. In the present case, the matrices B and V are given in equations (8) and (11), respectively. We thus find the following correlation matrix for η_w and η_z:

$$\tilde{B} = \begin{pmatrix} w + \frac{1}{2}w^2 - pz^2 & z(1 - 2\mu + w(\frac{1}{2} - p)) \\ z(1 - 2\mu + w(\frac{1}{2} - p)) & w + \frac{1}{2}w^2 - pz^2 \end{pmatrix}. \tag{14}$$

Notice that whilst the original noise variables η_x and η_y were independent (the off-diagonal entries of B were zero), the noise variables in the new coordinates are correlated with each other.

To enforce the assumed separation of time-scales between w and z, we impose the following conditions:

$$w = 1 + \sqrt{1 + 2pz^2} \quad \text{and} \quad \eta_w(t) \equiv 0\,. \tag{15}$$

The first of these sets w to its value on the nullcline for a given z, whilst the second removes the possibility of any noise-induced fluctuations. What effect do these constraints have on the evolution of z? First, we may substitute $w = 1 + \sqrt{1 + 2pz^2}$ into (12) to remove the dependence on w, thus

$$\frac{dz}{dt} = f(z) + \eta_z(t)\,, \tag{16}$$

where

$$f(z) = z\left(1 - 2\mu - \left(\frac{1}{2} - p\right)\left(1 + \sqrt{1 + 2pz^2}\right)\right). \tag{17}$$

Second, we must determine the effect of the conditions on the noise variable η_z. Since η_w and η_z are correlated, imposing $\eta_w = 0$ will alter the statistical distribution of η_w.

Again we apply a general result on correlated Gaussian random variables (see Appendix B of [16]). Suppose that a collection of Gaussian random variables $(\eta_1, \ldots, \eta_n)$ has correlation matrix $\tilde{B}$. Let $\bar{B}$ be the correlation matrix of $(\eta_2, \ldots, \eta_n)$ conditioned on the event that $\eta_1 = 0$. Then $\bar{B}$ and $\tilde{B}$ are related by

$$[\bar{B}^{-1}]_{ij} = [\tilde{B}^{-1}]_{ij}\,, \quad \text{for all} \ \ i, j = 2, \ldots, n\,. \tag{18}$$

This fact straightforward to prove by integration of the Gaussian probability density function. In particular, if $n = 2$ then the variance of η_2 conditioned on $\eta_1 = 0$ is

$$\bar{B} = \tilde{B}_{22} - \frac{\tilde{B}_{12}\tilde{B}_{21}}{\tilde{B}_{11}}\,. \tag{19}$$

Applying formula (19) to the correlation matrix found in (14), we obtain

$$\langle \eta_z(t)\eta_z(t') \rangle = \varepsilon\,\delta(t - t')g(z)\,, \tag{20}$$
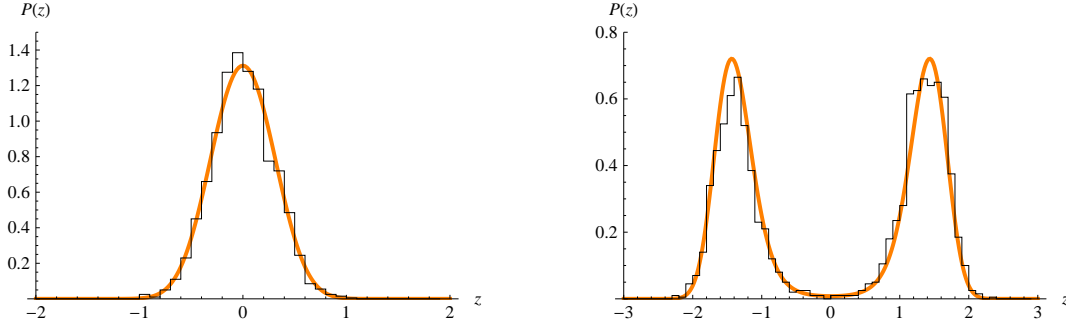
Figure 2. The stationary distribution of $z = x - y$ as predicted by the reduced-dimension model (orange curve) and measured from a single long simulation run of the individual-based model (black histogram). The parameters are the same as those in figure 1. The theoretical prediction was obtained by numerically integrating (22), with parameters taken from figure 1, whilst 2000 data points were collected from simulation run at time intervals of 100.

where the noise strength g is given by

$$g(z) = \left(w + \frac{1}{2}w^2 - pz^2 \right) \left(1 - \left(z \frac{1 - 2\mu + w(\frac{1}{2} - p)}{w + \frac{1}{2}w^2 - pz^2} \right)^2 \right) , \tag{21}$$

and where of course $w = 1 + \sqrt{1 + 2pz^2}$.

Equations (16) and (20) together define a one-dimensional stochastic differential equation. Although it may look a little complicated, being one-dimensional means that most questions of interest about this system can be answered by standard methods [12]. For example, the long-time average behaviour of the model is captured in the stationary distribution of z, which has the following explicit form:

$$P(z) = \frac{1}{g(z)} \exp \left(\frac{2}{\varepsilon} \int_{-\infty}^{z} \frac{f(z')}{g(z')} \, dz' \right) . \tag{22}$$

In figure 2 we compare the analytical prediction of (22) with a histogram of the z-coordinate of the sample points of stochastic simulations taken from figure 1. Clearly, the reduced one-dimensional model provides a very good fit to the data coming from the individual-based simulation. It should also be pointed out that although we have developed the theory based on the local behaviour around the coexistence fixed point $(1, 1)$, the approximation remains successful even in the unstable regime.

2.3. General formulation

We close this section by providing a description of the method for an arbitrary n-dimensional SDE

$$\frac{d\boldsymbol{x}}{dt} = \boldsymbol{A}(\boldsymbol{x}) + \boldsymbol{\eta}(t) , \tag{23}$$

with noise correlations

$$\left\langle \eta_i(t)\eta_j(t') \right\rangle = \varepsilon \, \delta(t - t') \left[B(\boldsymbol{x}) \right]_{ij} , \tag{24}$$

where $i, j = 1, \ldots n$.

Suppose we are interested in behaviour around a fixed point $\boldsymbol{x}^*$. Let J be the Jacobian of $\boldsymbol{A}$ at that point and write $\lambda_1, \ldots, \lambda_n$ for its eigenvalues. Suppose further that λ_1 is non-degenerate, real, large and negative, thus its associated eigenvector $\boldsymbol{v}_1$ represents a very stable direction. We aim to eliminate fluctuations in this direction to produce a reduced-dimension model. To apply our method, we first make a change of variables from the n-vector $\boldsymbol{x}$ to a single variable w and an $(n-1)$-vector $\boldsymbol{z}$, via the coordinate transformation

$$\begin{pmatrix} w \\ \boldsymbol{z} \end{pmatrix} = V\boldsymbol{x} \,. \tag{25}$$

It is possible to choose V so that

$$\bar{J} = VJV^{-1} = \begin{pmatrix} \lambda_1 & 0 \\ 0 & L \end{pmatrix}, \tag{26}$$

where L is an $(n-1) \times (n-1)$ matrix. The first row of V must be $\boldsymbol{v}_1^T$, and therefore near the fixed point $\boldsymbol{x}^*$, the variable w describes the distance from the slow manifold along the fast direction $\boldsymbol{v}_1$, while the remaining $n-1$ slow degrees of freedom are captured by $\boldsymbol{z} = (z_2, \ldots, z_n)^T$. In the new coordinates the SDE becomes

$$\frac{d}{dt}\begin{pmatrix} w \\ \boldsymbol{z} \end{pmatrix} = V\boldsymbol{A}\left(V^{-1}\begin{pmatrix} w \\ \boldsymbol{z} \end{pmatrix} \right) + \boldsymbol{\zeta}(t) \,, \tag{27}$$

where

$$\left\langle \zeta_i(t)\zeta_j(t') \right\rangle = \varepsilon\, \delta(t - t')\left[\tilde{B}(w, \boldsymbol{z}) \right]_{ij}, \tag{28}$$

and $\tilde{B} = VBV^T$. We wish to constrain w to its nullcline, which is defined by the equation

$$\boldsymbol{v}_1^T\boldsymbol{A}\left(V^{-1}\begin{pmatrix} w \\ \boldsymbol{z} \end{pmatrix} \right) = 0 \,. \tag{29}$$

We will assume that we may unambiguously describe the nullcline by a known function $w = \theta(\boldsymbol{z})$; in the illustrative example, $\theta(z)$ was given by (15), a branch of a hyperbola. In general however, it may have a more complicated form.

To enforce the assumed separation of time-scales between w and the other variables, we impose the following conditions:

$$w = \theta(\boldsymbol{z}) \quad \text{and} \quad \zeta_1(t) \equiv 0 \,. \tag{30}$$

For the remaining variables, we have

$$\frac{d\boldsymbol{z}}{dt} = \bar{\boldsymbol{A}}(\boldsymbol{z}) + \boldsymbol{\zeta}(t) \,, \tag{31}$$

where

$$\left\langle \zeta_i(t)\zeta_j(t') \right\rangle = \varepsilon\, \delta(t - t')\left[\bar{B}(\boldsymbol{z}) \right]_{ij}, \quad i, j = 2, \ldots n. \tag{32}$$

The drift vector $\bar{\boldsymbol{A}}(\boldsymbol{z})$ and diffusion matrix $\bar{B}(\boldsymbol{z})$ are derived from $\boldsymbol{A}$ and $\tilde{B}$ as follows. For $i, j = 2, \ldots, n$

$$\left[\bar{\boldsymbol{A}}(\boldsymbol{z})\right]_i = \left[V\boldsymbol{A}\left(V^{-1}\begin{pmatrix}\theta(\boldsymbol{z}) \\ \boldsymbol{z}\end{pmatrix}\right)\right]_i \tag{33}$$

and

$$\left[\bar{B}(\boldsymbol{z})^{-1}\right]_{ij} = \left[\tilde{B}(\theta(\boldsymbol{z}), \boldsymbol{z})^{-1}\right]_{ij}. \tag{34}$$

The general solution for the new noise covarience matrix can then be shown to be

$$\bar{B}(\boldsymbol{z}) = \mathcal{B}_{22}(\boldsymbol{z}) - \mathcal{B}_{21}(\boldsymbol{z})\mathcal{B}_{11}^{-1}(\boldsymbol{z})\mathcal{B}_{12}(\boldsymbol{z}), \tag{35}$$

where $\mathcal{B}$ are matrices of the partitioned $\tilde{B}$ matrix;

$$\tilde{B} = \begin{pmatrix} \mathcal{B}_{11}(\boldsymbol{z}) & \mathcal{B}_{12}(\boldsymbol{z}) \\ \mathcal{B}_{21}(\boldsymbol{z}) & \mathcal{B}_{22}(\boldsymbol{z}) \end{pmatrix}, \tag{36}$$

with $\mathcal{B}_{11}(\boldsymbol{z})$ a 1×1 matrix and $\mathcal{B}_{22}(\boldsymbol{z})$ an $(n-1) \times (n-1)$ matrix (see Appendix B of [16]). Equations (31) and (32) describe a reduced-dimension stochastic system in which the fast direction associated with the eigenvalue λ_1 has been eliminated.

3. Application: seasonally forced epidemics

3.1. Model definition and deterministic treatment

The SEIR model is a simplified epidemiological model describing the spread of a disease through a population [21]. Members of the population may be in one of four states: susceptible (S), exposed (E), infectious (I) and recovered (R). The susceptible individuals come into contact with the infected and become exposed with infection rate $\beta(t)$, which may vary with time. Those exposed to the disease then become infectious with a rate of disease onset α. Finally, the infectious recover with an average rate of γ. In addition to these disease dynamics, there is a constant birth and death rate μ; it is traditional to hold the population size constant by treating death and birth as a single process whereby an individual returns to the susceptible state. As in the earlier illustrative model, the dynamics may be conveniently summarized using the notation of chemical reactions:

$$\begin{array}{llll} \text{Infection:} & S + I \xrightarrow{\beta(t)} E + I & \text{Incubation:} & E \xrightarrow{\alpha} I \\ \text{Recovery:} & I \xrightarrow{\gamma} R & \text{Death/Birth:} & E, I, R \xrightarrow{\mu} S. \end{array} \tag{37}$$
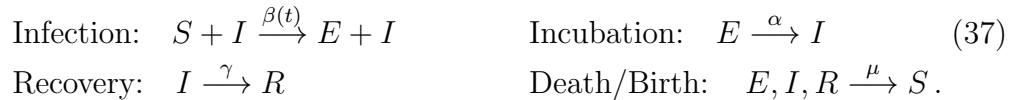
We write n_S, n_E, n_I, n_R for the number of individuals in states S, E, I and R, respectively. The total population size is then given by $N = n_S + n_E + n_I + n_R$, which does not vary, meaning that there are three degrees of freedom. With just a slight abuse of notation we introduce variables $S = n_S/N$, $E = n_E/N$ and $I = n_I/N$ which describe the population density of individuals in each disease state. Note that there is no need for a variable associated to the recovered state, since the conservation of total population

makes it a dependent variable. Applying the same system-size expansion as before (see Appendix A for details), we obtain the following effective SDE system:

$$\frac{dS}{dt} = \mu(1 - S) - \beta(t)SI + \eta_1(t) \,,$$
$$\frac{dE}{dt} = \beta(t)SI - (\mu + \alpha)E + \eta_2(t) \,, \tag{38}$$
$$\frac{dI}{dt} = \alpha E - (\mu + \gamma)I + \eta_3(t) \,,$$

where $\eta_{1,2,3}$ are Gaussian white noise variables with correlations

$$\left\langle \eta_i(t)\eta_j(t') \right\rangle = \frac{1}{N}\delta(t - t')B_{ij} \,,$$

$$B = \begin{pmatrix} \mu(1 - S) + \beta(t)SI & -\mu E + \beta(t)SI & -\mu I \\ -\mu E + \beta(t)SI & \beta(t)SI + (\mu + \alpha)E & -\alpha E \\ -\mu I & -\alpha E & \alpha E + (\mu + \gamma)I \end{pmatrix} \,. \tag{39}$$

In this section we discuss the behaviour of the model in the deterministic limit $N \to \infty$, before moving on to consider the full stochastic system in Section 3.2. When the infection rate is not seasonally forced, so $\beta(t) \equiv \beta$, there are a pair of fixed points. The first of these represents the extinction of the disease: $S = 1\,, E = 0\,, I = 0$. The second fixed point has coordinates

$$S^* = \frac{(\alpha + \mu)(\gamma + \mu)}{\alpha\beta} \,, \quad E^* = \frac{\mu(1 - S_*)}{\alpha + \mu} \,, \quad I^* = \frac{\alpha\mu(1 - S_*)}{(\alpha + \mu)(\gamma + \mu)} \,, \tag{40}$$

and is referred to as the endemic state. We are concerned with the regime in which the endemic state is stable and the extinct state is unstable, which holds for a range of epidemiologically realistic parameter values. In general, the rate parameter μ controlling birth and death will be much smaller than the remaining rate parameters β, α and γ, since this takes place on a much longer timescale than the disease dynamics. The parameter μ can therefore be utilized as an expansion parameter to simplify some of the expressions in the analysis. To first order in μ, the eigenvalues of the Jacobian at the endemic state are

$$\lambda_1 = -(\alpha + \gamma) - \mu\frac{\alpha(2\alpha + \beta) + 3\alpha\gamma + 2\gamma^2}{(\alpha + \gamma)^2} \,, \tag{41}$$

and the complex-conjugate pair

$$\lambda_{2,3} = \mu\frac{\alpha^2\beta + \beta\gamma^2 + \alpha\gamma(\beta + \gamma)}{2\gamma(\alpha + \gamma)^2} \pm i\sqrt{\mu\frac{\alpha(\beta - \gamma)}{\alpha + \gamma}} \,. \tag{42}$$

Notice that we have a separation of time-scales: $\text{Re}[\lambda_1] \ll \text{Re}[\lambda_{2,3}]$, meaning that λ_1 corresponds to a highly stable direction. This behaviour has been previously noted and exploited in the deterministic setting [22]. In addition, since μ is small, the imaginary parts of $\lambda_{2,3}$ are typically larger than the real parts, meaning that we may expect highly oscillatory trajectories in the neighbourhood of the endemic state. We can thus expect
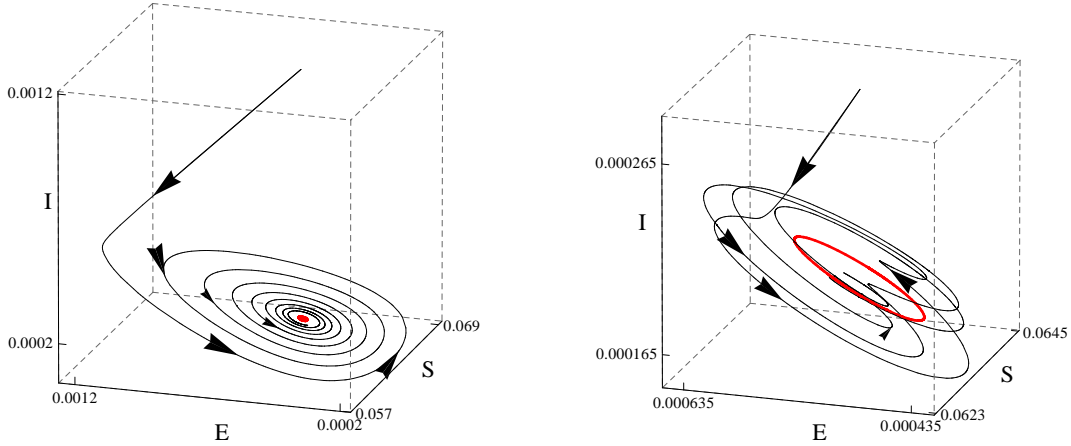
Figure 3. Deterministic trajectories for the SEIR model. Left: No seasonal forcing ($\delta = 0$), fixed point shown in red. Right: system in the presence of forcing, ($\delta = 0.02$), limit cycle highlighted in red. Remaining parameters in both plots are $\beta_0 = 1575$ year^{-1}, $\alpha = 35.84$ year^{-1}, $\gamma = 100$ year^{-1} and $\mu = 0.02$ year^{-1} which are standard choices for measles [19, 23].

the system to first collapse rapidly in the direction of the first eigenvector, followed by a slow, almost-planar, spiralling decay to the endemic state, as shown in figure 3.

Introducing seasonal forcing of the infection rate creates an additional layer of complexity. A typical choice would be

$$\beta(t) = \beta_0 \Big(1 + \delta \cos(2\pi t) \Big), \tag{43}$$

where β_0 describes the basal infection rate, δ is the forcing amplitude, and time is measured in years. The deterministic system will now not settle to the endemic state, but instead exhibit limit cycle behaviour. For the sake of simplicity, we will consider only parameter values which result in a single stable limit cycle of period $T = 1$ year.

Similar to linear stability analysis of fixed points, there is a well-developed theory for analysing perturbations around limit cycles. Writing $(S^*(t), E^*(t), I^*(t))$ for the limit cycle, we introduce the vector

$$\boldsymbol{\xi}(t) = \sqrt{N} \left(\begin{array}{c} S(t) - S^*(t) \\ E(t) - E^*(t) \\ I(t) - I^*(t) \end{array} \right), \tag{44}$$

where the factor $\sqrt{N}$ has been introduced for notational consistency with later sections. Since we are effectively linearizing the system, it should be noted that this factor has no effect on the following analysis. To first order then, the dynamics of $\boldsymbol{\xi}$ are governed by

$$\frac{d\boldsymbol{\xi}}{dt} = J(t)\boldsymbol{\xi}, \tag{45}$$

where J is the time-dependent Jacobian of (38). This equation may be solved using Floquet theory [24]. The key result of the theory states that solution trajectories

may be decomposed into the product of a periodic vector with an exponentially growing/decaying amplitude. General solutions are then of the form

$$\boldsymbol{\xi}(t) = \sum_{i=1}^{n} c_i \boldsymbol{q}_i(t) e^{\sigma_i t}, \tag{46}$$

where n is the number of degrees of freedom in the model, c_i is a constant, $\boldsymbol{q}_i(t)$ a periodic vector with the same period T as the limit cycle, and the value σ_i determining the rate of growth/decay is referred to as the Floquet exponent. Akin to the eigenvalues of the fixed point Jacobian, Floquet exponents are indicative of the stability of the limit cycle in the time-varying directions $\boldsymbol{q}_i(t)$; perturbations to the trajectory will grow if $\mathrm{Re}[\sigma_i] > 0$ and decay if $\mathrm{Re}[\sigma_i] < 0$.

In all but the most trivial examples, obtaining the Floquet exponents and periodic vectors must be carried out numerically. The procedure (detailed in Appendix B) requires first computing a matrix whose columns are the independent solutions to (45), $X(t) = (\boldsymbol{\xi}_1(t) \cdots \boldsymbol{\xi}_n(t))$, and then determining the eigenvalues of $X(T)$. If there is a rapid collapse along a stable direction, followed by slow dynamics in a subspace, then the columns of $X(T)$ will be almost linearly dependent, since all solution trajectories quickly move towards the subspace. In turn, the real parts of eigenvalues of the matrix $X(T)$ will differ by many orders of magnitude. Obtaining these eigenvalues accurately is crucial to the remaining analysis. However, if the disparity between the eigenvalues is too great, numerical procedures may not maintain sufficient accuracy in calculations involving the eigenvalues. This problem was previously highlighted in [19] within an analysis of the seasonally forced SEIR model. There, using the same epidemiological parameters as used in figure 3, the eigenvalues of $X(T)$ differed by an order of 10^{59}. This prompted the authors to implement arbitrary precision numerical methods, at a considerable cost in computing time. We propose that the general method detailed in Section 2.3 can be used as a technique to remove this fast direction and hence circumnavigate the numerical difficulties encountered in the analysis.

3.2. Stochastic treatment exploiting the slow manifold

Beginning with the unforced case, we let λ_1, λ_2 and λ_3 be as in (41, 42) and write $\boldsymbol{v}_1$, $\boldsymbol{v}_2$, $\boldsymbol{v}_3$ for the corresponding eigenvectors. We introduce the transformation matrix $V = (\boldsymbol{v}_1 \quad (\boldsymbol{v}_2 + \boldsymbol{v}_3) \quad i(\boldsymbol{v}_2 - \boldsymbol{v}_3))^T$ and new variables

$$\begin{pmatrix} w \\ z_2 \\ z_3 \end{pmatrix} = V \begin{pmatrix} S \\ E \\ I \end{pmatrix}. \tag{47}$$

The Jacobian of the transformed system at the endemic fixed point takes the form

$$\bar{J} = \begin{pmatrix} \lambda_1 & 0 \\ 0 & L \end{pmatrix} + \mathcal{O}(\mu^{3/2}), \quad \text{where} \quad L = \begin{pmatrix} \mathrm{Re}[\lambda_2] & \mathrm{Im}[\lambda_3] \\ \mathrm{Im}[\lambda_2] & \mathrm{Re}[\lambda_3] \end{pmatrix}. \tag{48}$$

The nullcline for w is determined by manipulating (29) into the form $w = \theta(z_2, z_3)$, with $\boldsymbol{A}$ copied from (38). Although an explicit form for θ can be found, the expression is too complicated to be worth reproducing here.

To capture the effects of stochasticity, we introduce variables describing the fluctuations in the new coordinates, rescaled by a factor of $\sqrt{N}$,

$$\bar{\boldsymbol{\xi}} = \sqrt{N} \begin{pmatrix} w - w^* \\ \boldsymbol{z} - \boldsymbol{z}^* \end{pmatrix}. \tag{49}$$

Making this substitution in (38) and keeping only first order terms in $1/N$ and μ, we find that $\bar{\boldsymbol{\xi}}$ obeys

$$\frac{d\bar{\boldsymbol{\xi}}}{dt} = \bar{J}\bar{\boldsymbol{\xi}} + \boldsymbol{\zeta}(t), \tag{50}$$

where

$$\langle \zeta_i(t)\zeta_j(t') \rangle = \delta(t - t')\tilde{B}_{ij}, \quad i, j = 1, 2, 3. \tag{51}$$

The matrix $\tilde{B}$ is given by

$$\tilde{B} = VBV^T \Big|_{(S,E,I)=(S^*,E^*,I^*)}, \tag{52}$$

where B is as in (39). Note that applying the constraint $w = \theta(z_2, z_3)$ induces the relationship

$$w^* + \frac{\bar{\xi}_1}{\sqrt{N}} = \theta\left(z_2 + \frac{\bar{\xi}_2}{\sqrt{N}}, z_3 + \frac{\bar{\xi}_3}{\sqrt{N}}\right). \tag{53}$$

Expanding once more in large N, we find

$$\bar{\xi}_1 = \left[\bar{\xi}_2 \frac{\partial\theta}{\partial z_2} + \bar{\xi}_3 \frac{\partial\theta}{\partial z_3}\right]\Big|_{\boldsymbol{z}=\boldsymbol{z}^*}. \tag{54}$$

After elimination of the fast direction, (50) becomes

$$\frac{d}{dt}\begin{pmatrix} \bar{\xi}_2 \\ \bar{\xi}_3 \end{pmatrix} = L\begin{pmatrix} \bar{\xi}_2 \\ \bar{\xi}_3 \end{pmatrix} + \begin{pmatrix} \zeta_2(t) \\ \zeta_3(t) \end{pmatrix}, \tag{55}$$

where ζ_2 and ζ_3 now have correlation matrix $\bar{B}$, which is related to $\tilde{B}$ by (34).

We move on now to study the situation of seasonally forced infection rate. In principle, the calculations above apply only in the limit of small forcing amplitude (that is, $\delta \to 0$ in (43)). We learnt in the illustrative ecological example, however, that although our theory is developed to apply in the locality of a stable fixed point, it can continue to provide a useful approximation if this condition does not strictly hold. Applying that lesson to the present case, we modify (55) to allow L and $\bar{B}$ to become functions of time, as dictated by the replacement $\beta \mapsto \beta(t)$. Essentially, we are approximating the limit cycle by its projection on to the nullcline of the fast direction at the endemic fixed point of the unforced model, and then demanding that any stochastic fluctuations remain on this nullcline. The application of such a coordinate change to the forced system can be further motivated by a consideration of the periodic matrix $J(t)$ in (45) in the limit of small forcing.

A Floquet analysis of the deterministic part of the reduced system finds two complex conjugate Floquet multipliers, with the third disparate multiplier having been eliminated from the system. This system no longer suffers from the numerical difficulties which plague the full three-dimensional model.

To quantify the effect of stochastic fluctuations in this model, we follow the standard procedure of computing the autocorrelation matrix $\mathcal{C}(\tau)$ of oscillations around the limit cycle, which has entries

$$\left[\mathcal{C}(\tau)\right]_{ij} = \frac{1}{T}\int_0^T \left\langle \bar{\xi}_i(t)\bar{\xi}_j(t+\tau)\right\rangle dt\,. \tag{56}$$

Of course our reduced system (55) is two-dimensional, meaning that the entries of $\mathcal{C}$ pertaining to $\bar{\xi}_1$ must be deduced from (54). The coordinate transformation applied at the start in (47) may then be inverted to give the autocorrelation matrix for the fluctuations in S, E and I.

The Fourier transform of the diagonal entries of the autocorrelation gives the power spectrum of oscillations, which provides a convenient visualization of the stochastic fluctuations. In figure 4 we plot the power spectrum of fluctuations in the number of infected individuals around the limit cycle, comparing between stochastic simulations and the theoretical prediction using the reduced-dimension model (55). The agreement between the simulation and theory can be seen to be excellent; the spectra lie virtually coincident. The peaks in the theoretical spectrum are found at the same positions as those for the simulated spectrum. These are approximately given by

$$v_j = \frac{j}{T} \pm \frac{|\mathrm{Im}[\sigma_{1,2}]|}{2\pi}\,, \tag{57}$$

where j is an integer and σ_i are the Floquet exponents as defined in Appendix B. This is in agreement with [19, 25] where peaks at these positions were also found. The overall benefit that was garnered by the procedure however was that of computational efficiency. The computation of the theoretical power spectra approximated from the reduced system takes only a small fraction of the computing time of the full system.

4. Conclusion

In this article we have introduced a systematic and general procedure to eliminate fast degrees of freedom in stochastic dynamical systems. The method was inspired by the highly successful theory of slow manifolds in deterministic systems, from which we borrow the basic notion of restricting attention to trajectories occupying a low-dimensional subspace. In the stochastic setting, we achieve this by enforcing the condition that the system state remains fixed to the nullcline of the fast direction. We apply this condition to SDE systems with correlated Gaussian white noise, obtaining a lower-dimensional system in which the fast degree of freedom has been eliminated. Importantly, the reduced system has the same basic form as the original, meaning that we have not inadvertently introduced any complicating factors such as non-Markovian processes or coloured noise.
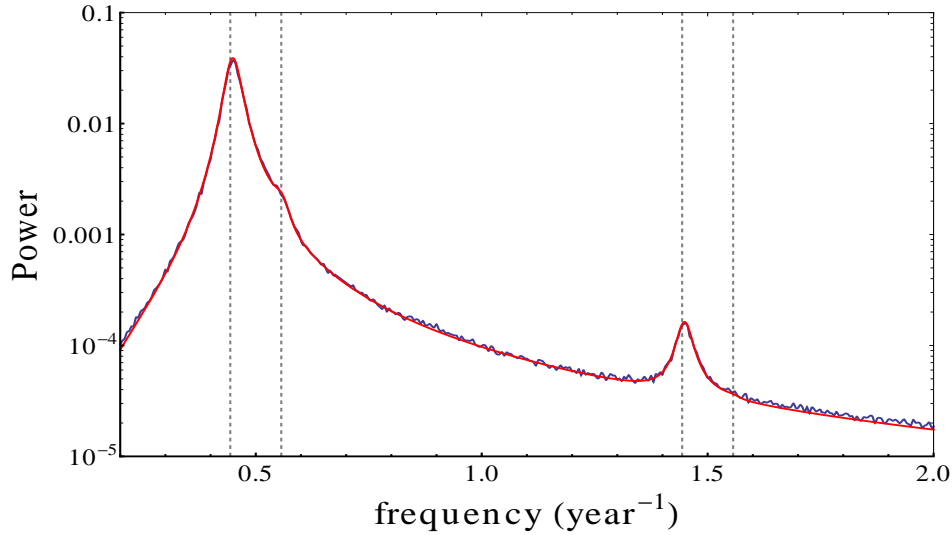
Figure 4. Power spectra for the fluctuations in the number of infected about the limit cycle. The analytical power spectrum calculated using the slow manifold approximation is plotted in red while the power spectrum from stochastic simulations is in blue. Agreement is such that the spectra are difficult to distinguish; the spectrum from simulated results is primarily discernible through its stochastic nature relative to the smooth analytical line. Dotted lines indicate the position of the peaks in the power spectra given by (57). Epidemiological parameters are $\beta_0 = 1575$ year^{-1}, $\alpha = 35.84$ year^{-1}, $\gamma = 100$ year^{-1} and $\mu = 0.02$ year^{-1}. The simulated spectrum is calculated as the average power spectrum of 1000 stochastic realizations, each lasting 200 years with a system size of $N = 10^8$.

Although our procedure is by no means the only way to reduce the dimension of a stochastic dynamical system, it does offer some distinct advantages. We would argue that our approach is relatively simple: working in the SDE setting makes clear the analogy to deterministic slow manifold theory, and our basic idea (confining the stochastic system to the slow manifold) is intuitively easy to grasp. Moreover, the technique is generally applicable; it does not require an explicitly fast variable, since the fast direction is determined instead from a linear stability analysis. Nor does the technique itself require knowledge of the deterministic trajectories, as is necessary in [15]. Finally, as we have shown, the approximation often remains successful even outside of the parameter regime in which it is developed. Confinement to the slow manifold does, however, remove any effects arising from the nature of the flow field away from the manifold. For example, in [15] it was shown that the way in which stochastically perturbed trajectories relax back can induce a small flow along the manifold; this effect is overlooked by our method. Looking to the future, there is the possibility to developing a general framework which includes effects of this type, and we hope that the work presented in this article will provide a significant step towards that goal.

Acknowledgments

Appendix A.

In this appendix we give the details of the derivation of the SDE systems (7) and (38), which are mesoscopic descriptions of the individual based models used as examples in sections 2 and 3, respectively. The general formulation of the method is described in [17].

Recall the model studied in section 2, with reactions listed in (6). The state of the system at a given time is specified by the vector $\boldsymbol{n} = (n_X, n_Y)$, where n_X and n_Y denote the number of individuals in populations X and Y, respectively. We write $P(\boldsymbol{n}, t)$ for the probability that the model is in state $\boldsymbol{n}$ at time t. To properly specify the model, we must write a master equation describing the time evolution of P.

The reactions (6) define the possible ways in which the system can transition from the state $\boldsymbol{n}'$ to the state $\boldsymbol{n}$; the rate with which these transitions occur is written $T(\boldsymbol{n}|\boldsymbol{n}')$. Specifically, we have

$$
\begin{aligned}
T(n_X + 1, n_Y | n_X, n_Y) &= (1 - \mu)n_X + \mu n_Y \\
T(n_X, n_Y + 1 | n_X, n_Y) &= \mu n_X + (1 - \mu)n_Y \\
T(n_X - 1, n_Y | n_X, n_Y) &= \varepsilon(1/2 - p)n_X^2 + \varepsilon(1/2 + p)n_X n_Y \\
T(n_X, n_Y - 1 | n_X, n_Y) &= \varepsilon(1/2 + p)n_X n_Y + \varepsilon(1/2 - p)n_Y^2 \, .
\end{aligned}
\tag{A.1}
$$

The master equation is written in terms of the transition rates in the following way;

$$
\frac{\partial P(\boldsymbol{n}, t)}{\partial t} = \sum_n [T(\boldsymbol{n}|\boldsymbol{n}')(t)P(\boldsymbol{n}', t) - T(\boldsymbol{n}'|\boldsymbol{n})(t)P(\boldsymbol{n}, t)] \, .
\tag{A.2}
$$

While the master equation describes the dynamics of the system entirely, it is in general not analytically tractable. A particular realization of a process obeying the master equation may be simulated using the Gillespie algorithm [20]. While the simulation does not necessarily provide as deep an understanding of the system as an analytical treatment, it provides a useful comparison for the analytical techniques.

To make analytical progress, we consider the limit of small ε, applying a procedure known as the Kramers-Moyal expansion [17, 18]. Briefly, it involves the introduction of the rescaled state vector $\boldsymbol{x} = \varepsilon \boldsymbol{n}$, followed by a Taylor expansion of T and P around $\boldsymbol{x}$. Truncating after two terms, results in a Fokker-Planck equation [26] describing the evolution of the probability density function;

$$
\frac{\partial P(\boldsymbol{x}, t)}{\partial t} = -\sum_i \frac{\partial}{\partial x_i} [A_i(\boldsymbol{x})P(\boldsymbol{x}, t)] + \frac{\varepsilon}{2} \sum_{i,j} \frac{\partial^2}{\partial x_i \partial x_j} [B_{ij}(\boldsymbol{x})P(\boldsymbol{x}, t)] \, .
\tag{A.3}
$$

The vector $\boldsymbol{A}(\boldsymbol{x})$ and the matrix $B(\boldsymbol{x})$ can be calculated from the probability transition rates (A.1). It can be shown [12] that the above equation is equivalent to the SDE

$$\frac{d\boldsymbol{x}}{dt} = \boldsymbol{A}(\boldsymbol{x}) + \boldsymbol{\eta}(t), \tag{A.4}$$

defined in the sense of Itō [27]. Here $\boldsymbol{\eta}(t)$ is as usual a Gaussian white noise term such that $\langle \boldsymbol{\eta}(t) \rangle = 0$ and $\langle \eta_i(t)\eta_j(t') \rangle = \varepsilon B_{ij}\delta(t - t')$. For the present example, the SDE formulation is given in (7).

The procedure is much the same for the epidemiological model considered in section 4. The transition rates are this time derived from the reactions (37), and the state vector is now $\boldsymbol{n} = (n_S, n_E, n_I)$. As is usual in this type of model, we choose $1/N$ (where N is the total number of individuals in the population) for the small parameter, ϵ, used in the Kramers-Moyal expansion. (38) is found as the result.

Appendix B.

The analogue of a linear stability analysis for systems with periodic components is known as Floquet theory [24]. It can also play an important role in the analysis of stochastic fluctuations about a deterministic trajectory [19, 25]. In this appendix the general formulation of Floquet theory is discussed before the more detailed application to linear stochastic systems is given.

Floquet theory gives the solutions to sets of linear differential equations in the form of (45), where $J(t)$ is periodic with a period T. The general solution can be shown to be

$$\boldsymbol{\xi}(t) = \sum_{i=1}^{n} c_i \boldsymbol{q}_i(t) e^{\sigma_i t}, \tag{B.1}$$

where $\boldsymbol{q}_i(t)$ is a periodic vector and σ_i are termed the Floquet exponents of the system. The quantities $\rho_i = e^{\sigma_i T}$ are termed the Floquet multipliers of the system.

In particular one can work in a canonical form for calculational ease, with canonical quantities denoted with a superscript 0. The canonical form is constructed from n decomposed solutions to (45) such that $\boldsymbol{\xi}_i^{(0)}(t) = \boldsymbol{q}_i^{(0)}(t)e^{\sigma_i t}$. A fundamental matrix of these solutions may then be introduced along with matrices $Y^{(0)}$ and $Q^{(0)}$. For the case $n = 3$ these may be expressed as

$$X^{(0)} = [\boldsymbol{\xi}_1^{(0)}(t), \boldsymbol{\xi}_2^{(0)}(t), \boldsymbol{\xi}_3^{(0)}(t)], \tag{B.2}$$

$$X^{(0)} = Q^{(0)}Y^{(0)}, \tag{B.3}$$

$$Q^{(0)} = [\boldsymbol{q}_1^{(0)}(t), \boldsymbol{q}_2^{(0)}(t), \boldsymbol{q}_3^{(0)}(t)], \tag{B.4}$$

$$Y^{(0)} = Diag[e^{\mu_i t}]. \tag{B.5}$$

A method for obtaining the Floquet multipliers μ_i along with the canonical form of the solutions is now required. Obtaining both is dependent on the determination of a matrix known as the monodromy matrix, which we shall now discuss.

The monodromy matrix, D, is defined such that $X(t + T) = X(t)D$, for any fundamental matrix $X(t)$ constructed from linearly independent solutions to (45). It can be shown that while the monodromy matrix is dependent on the fundamental matrix chosen, its eigenvalues are not [24]. The eigenvalues of D are ρ_i, the Floquet multipliers of the system. Further, if a matrix W is constructed from the eigenvectors of D, the canonical fundamental matrix $X^{(0)}(t)$ is related to a general fundamental matrix $X(t)$ via $X^{(0)}(t) = X(t)W$. Therefore, the monodromy matrix allows the canonical fundamental matrix $X^{(0)}(t)$ to be determined from a general fundamental matrix $X(t)$, along with the matrix $Y^{(0)}$. From these the periodic matrix $Q^{(0)}(t)$ can also be deduced.

If the system under consideration displays limit cycle behaviour, the preceding analysis must evidently be carried out numerically. In general a fundamental matrix obtained numerically will have to be transformed into canonical form by a numerical determination of the monodromy matrix, $D = X^{-1}(t)X(t + T)$. For a system with initial conditions $t = 0$, $X(0) = I$, this simplifies to $D = X(T)$.

Now the stochastic system can be considered;

$$\frac{d\boldsymbol{\xi}}{dt} = J(t)\boldsymbol{\xi} + \boldsymbol{\eta}(t), \tag{B.6}$$

where $\boldsymbol{\eta}(t)$ is a vector of Gaussian white noise terms defined as in Appendix A, except that now the noise covariance matrix depends explicitly on time through the varying parameter $\beta(t)$; $\langle \eta_i(t)\eta_j(t') \rangle = \varepsilon B_{ij}(t)\delta(t - t')$. The solution may be constructed as a sum of the general solution to (45) along with a particular solution, so that

$$\boldsymbol{\xi}(t) = X^{(0)}(t)\boldsymbol{\xi}^{(0)} + X^{(0)}(t) \int_{t_0}^{t} \left[X^{(0)}(s)\right]^{-1} \boldsymbol{\eta}(s)ds, \tag{B.7}$$

or, setting the initial conditions in the infinite past and making a change of integration variable $s \to s' = t - s$

$$\boldsymbol{\xi}(t) = Q^{(0)}(t) \int_{t_0}^{t} Y^{(0)}(s') \left[Q^{(0)}(t - s')\right]^{-1} \boldsymbol{\eta}(t - s')ds'. \tag{B.8}$$

In the course of the analysis conducted in Section 3, $\boldsymbol{\xi}(t)$ represents some stochastic fluctuation around limit cycle behaviour. An obvious quantity of relevance is the power spectrum of such fluctuations. To obtain the power spectrum, one first calculates the two-time correlation function $C(t+\tau, t) = \langle \boldsymbol{\xi}(t+\tau)\boldsymbol{\xi}^T(t) \rangle$; substituting (B.8) one obtains

$$C_{ij}(t + \tau, t) = Q^{(0)}(t + \tau)Y^{(0)}(\tau)\Lambda(t) \left[Q^{(0)}(t)\right]^T, \tag{B.9}$$

with

$$\Lambda(t) = \int_{t_0}^{\infty} Y^{(0)}(s)\Gamma(t - s)Y^{(0)}(s)ds \tag{B.10}$$

and

$$\Gamma(s) = \left[Q^{(0)}(s)\right]^{-1} B(s) \left[\left[Q^{(0)}(s)\right]^{-1}\right]^T. \tag{B.11}$$

The correlation function, $\mathcal{C}(\tau)$ is then simply related to the two-time correlation function by

$$\mathcal{C}(\tau) = \frac{1}{T}\int\limits_0^T C(t+\tau,t)dt. \tag{B.12}$$

In turn, the Weiner-Khinchin theorem tells us that the power spectrum, $P(\omega)$, is simply the Fourier transform of the correlation function, and so

$$P_i(\omega) = \int \mathcal{C}_{ii}(\tau)e^{i\omega\tau}d\tau. \tag{B.13}$$

The intermediate steps are left to the reader, but full details are found in [28]. A key point to note is that Equations (B.7-B.11) hold *only* for the canonical matrices $X^{(0)}$, $Q^{(0)}$ and $Y^{(0)}$.

References

[1] Wiggins S 2003 *Introduction to Applied Nonlinear Dynamical Systems and Chaos* (New York: Springer)
[2] Schöner G and Haken H 1986 *Z. Phys. B.* **63** 493–504
[3] Schöner G and Haken H 1987 *Z. Phys. B.* **68** 89–103
[4] Gardiner C W 1984 *Phys. Rev. A* **29** 2814–2822
[5] Titulaer U M 1983 *Z. Phys. B.* **50** 71–77
[6] Thomas P, Grima R and Straube A V 2012 *Phys. Rev. E.* **86** 041110
[7] Serra R, Andretta M, Compiani M and Zanarini G 1986 *Introduction to the Physics of Complex Systems* (Oxford: Pergamon Press)
[8] Contou-Carrere M N, Sotiropoulos V, Kaznessis Y N and Daoutidis P 2011 *Systems and Control Letters* **60** 75–86
[9] Wang W and Roberts A J 2013 *Journal of Mathematical Analysis and Applications* **398** 822–839
[10] Lan Y, Elston T C and Papoian G A 2008 *J. Chem. Phys.* **129** 214115
[11] Rogers T, McKane A J and Rossberg A G 2012 *Europhys. Lett.* **97** 40008
[12] Gardiner C W 2009 *Handbook of Stochastic Methods* (Berlin: Springer)
[13] Itō K 1951 *Mem. Am. Math. Soc.* **4** 1–51
[14] Haken H and Wunderlin A 1982 *Z. Phys. B.* **47** 179–187
[15] Lin Y T, Kim H and Doering C R 2012 *J. Stat. Phys.* **148** 646–662
[16] Rogers T, McKane A J and Rossberg A G 2012 *Phys. Biol.* **9** 066002
[17] McKane A J, Biancalani T and Rogers T 2012 *arXiv: 1211.0462*
[18] van Kampen N G 2007 *Stochastic Processes in Physics and Chemistry* (Amsterdam: Elsevier)
[19] Rozhnova G and Nunes A 2010 *Phys. Rev. E* **82** 041906
[20] Gillespie D T 1976 *J. Comp. Phys.* **22** 403–434
[21] Keeling M J and Rohani P 2007 *Modelling Infectious Diseases in Humans and Animals* (Princeton: Princeton University Press)
[22] Schwartz I B and Smith H L 1983 *J. Math. Biol.* **18** 233–253
[23] Schwartz I B 1985 *J. Math. Bio.* **21** 347–361
[24] Grimshaw R 1990 *Nonlinear Ordinary Differential Equations* (Oxford: CRC Press)
[25] Black A J and McKane A J 2010 *J. Theor. Biol.* **267** 85–94
[26] Risken H 1989 *The Fokker-Planck Equation* (Berlin: Springer)
[27] Jacobs K 2010 *Stochastic Processes for Physicists* (Cambridge: Cambridge University Press)
[28] Boland R P, Galla T and McKane A J 2009 *Phys. Rev. E* **79** 051131